

IN THE SPECIFICATION

Please amend paragraph [0025] of the Specification as follows:

What follows makes reference to a dynamic session management of virtualized network object IDs such as servers $s_1 \dots s_4$, referred to herein as “S1 . . . S4;” clients $c_1 \dots c_8$, referred to herein as “C1 . . . C8;” the elements of the session management model are: the network 3D maps $s_1(s_{11}, s_{12}, s_{13}), \dots, c_1(C_{11}, c_{12}, c_{13}), \dots$, referred to herein as “S1(S11, S12, S13), . . . C1(C11, C12, C13), . . .” containing the position of the servers and clients in the network space; the application/virtual space partitioning partitions, referred to herein as “P1 . . . P16;” the client partition mapping Client-Partition map describing the interest of the clients in application space partitions $c_1(P_1, P_2); c_2(P_6, P_7, P_8); c_3(P_1, P_2, P_5)$, referred to herein as: “C1(P1, P2); C2(P6, P7, P8); C3(P1, P2, P5);” the content objects, referred to herein as: “o1 . . . o10;” the content objects partitioning/replication map, referred to herein as: “P1(o1, o2, o3); P2(o4);” and the Server-Partition map $s_1(P_1, P_2, P_5), s_2(P_1, P_2, P_4)$, referred to herein as “S1 (P1, P2, P5), S2(P1, P2, P4).” The maps are used to construct a multi-attribute indexing structure that indexes the network and application level objects; the indexed network and application objects are grouped such that to satisfy network/application constraints as response time (e.g. server S1 s_1 satisfies the grouping constraint--the network delay to clients C1, C3 c_1, c_3 is bounded), capacity constraints. This result of the grouping process are membership lists containing network objects: e.g. $(s_1)(c_1, c_3)$ (S1)(C1, C3), the partitions associated with each group: (P1, P2) and the application objects corresponding to these partitions (o1, o2, o3, o4). The reference signs used herein to refer to the various object IDS are used for the purposes of clarifying explanation. Some of these reference signs may be used further in the Figures that follow.

Please amend Paragraph [0027] of the Specification as follows:

While FIG. 2 illustrates how network resources can be mapped to achieve specific communication goals, the principles of the present invention also manage communications based on the application space partitioning. For example, FIG. 3

illustrates an application space 300 that is divided based on communication interests into 16 partitions. Partition P1 (element 302) and partition P2 (element 304) are of communication interest to a client C1, as indicated by box 308 drawn around partition P1 and partition P2, while partition P5 (element 306) is not of interest to client C1. Partitions P1, P2, and P5 are mapped to application server S1. Furthermore, the communication interest represented by partition P1 include the content object o1. Thus, an application can be characterized by an application space 300 that can be partitioned into areas of interest to some clients, but not others, a given client may have interests in multiple partitions, and at least some of the partitions are made available to network overlay nodes.

Please amend Paragraph [0030] of the Specification as follows:

Yet a more general representation of a map is an application object map in which objects (communication interest objects, content objects, etc.) defined by an application are partitioned ~~are partitioned~~ in a space defined by a set of associated descriptors. The objects ~~and~~ are then mapped to network resources. The mapping of application objects partitioning into network resources can be done by allocating multiple partitions per network resource and replication of a partitions at multiple network resources.

Please amend Paragraph [0031] of the Specification as follows:

The use of maps in a communication session and data distribution management system 400 is illustrated in Figure 4. As shown, communication interest partitions 402, content partitions 404, and virtualized network objects 406 are indexed to form a multi-attribute space 408. Based on that indexing group membership is managed 410. This is performed by clustering users/clients according to their communication interests and network objects attributes as indexed in the multi-attribute space 408. The result is a step 412 of forming a membership list that assigns users to middleware communication nodes and an application index that assigns partitions to middleware communication nodes. Based on the membership lists and the application index, at step 414 the communications are managed using a hierarchical structure that is based on the

application index and membership lists. This enables multicast communications by having application servers send specific communications to specific middleware communication nodes that have users that are interested in such communications, and by having those middleware communication nodes forward communications only to the users that are interested in such communications. The application servers know which middleware communication nodes to send what information to, while the middleware communication nodes know who is to receive that information.

Please amend Paragraph [0046] of the Specification as follows:

As an example, FIG. 7 illustrates hierarchical group management using communication interest aggregation with proximity-based clustering of client nodes and interest aggregation in the hierarchical control topology. ~~Control~~ Server node s1 acts as a control node which indexes the communication interest of clients {c1,c2,c3,c4} and propagates control messages containing descriptors for active groups in the control hierarchy--e.g., when client node c sends a join message indicating its communication interest, s1 inserts a new group with descriptor SD1, adds client c, and propagates the control message to s4 in the control hierarchy. The control node s1 indexes the membership lists: ML1(SD1)={c1} and membership list ML2(SD2)={c2,c3,c4,s2,s3}; The ~~control server~~ node s3 acts as a control node which indexes the lists: ML(SD3)={c6, c7} and ML4(sd2)={C5,S1}. The control message to s4 cause s4 to add s1 to the group for SD1 (inserting a new group for SD1 in case there were no previous members with such interest) and propagates control messages further in the hierarchy, thus aggregating the communication interest of the lower level nodes. S1 also maps the membership lists corresponding to group descriptors SD1 and SD2 into multicast groups. In FIG. 8 the communication interest is indexed in a hierarchical structure with the membership lists corresponding to communication groups stored at the leaf level. The hierarchical structure includes interior nodes [D1, D2] and [D3,D4], as well as leaf nodes [d1,d2], [d3,d4], and [D5,D6]. Membership lists (e.g., [d1] x [d2]; [d3] x [d4]; [d5] x [d6]; [D1] x [D2]; and [D3] x [D4]) indexed at different nodes may refer to the same descriptor; since the data control uses the session descriptors, the membership lists can be local to

control nodes. The communication interest index is queried at control nodes to retrieve membership lists composition; e.g. the query on SD2 at node s4 retrieves {s1,s2,s3}, while a query in SD2 at s1 retrieves client list {c2,c3,c4} and forwarding server list {s2,s3}.

Please amend Paragraph [0061] of the Specification as follows:

Receiver/Sender Join/Leave Node Join: